IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application No.: 10/624,396	§	Examiner: Sorrell, Eron. J.
Filed: July 22, 2003	§	Group/Art Unit: 2182
Inventor(s):	§	Atty. Dkt. No: 5760-12100
Agarwal, et al.	§	Confirm No. 6815
	§	
	§	****CERTIFICATE OF E-FILING TRANSMISSION****
	§	I hereby certify that this correspondence is being transmitted via electronic filing to the United States Patent and Trademark Office on the date shown below
Title: TECHNIQUE TO MONITOR APPLICATION BEHAVIOR AND TUNE REPLICATION PERFORMANCE	§	Rory D. Rankin (Printed Name)
	§	/ Rory D. Rankin / (Signature) August 6, 2007 (Date)
	§	

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal of March 27, 2007, and Panel Decision dated June 6, 2007, Appellants present this Appeal Brief. Appellants respectfully request that this appeal be considered by the Board of Patent Appeals and Interferences.

I. REAL PARTY IN INTEREST

The subject application is owned by Veritas Operating Corporation, a corporation organized and existing under and by virtue of the laws of the State of Delaware, and having its principal place of business at 350 Ellis Street, Mountain View, CA, 94043, as evidenced by the assignment recorded at Reel 014336, Frame 0600.

II. RELATED APPEALS AND INTERFERENCES

No other appeals or interferences are known which would directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-2, 4, 6-10, 12, 14-18, and 20-23 are pending and rejected, and are the subject of this appeal. A copy of claims 1-2, 4, 6-10, 12, 14-18, and 20-23 as on appeal is included in the Claims Appendix hereto.

IV. STATUS OF AMENDMEMNTS

No amendments to the claims have been filed subsequent to the final rejection.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

The subject matter of the present claims generally relates to the field of computing systems and dynamic modification of system resources.

Claim 1 recites a computing system including an application, a storage device, a memory pool, and a replicator component. (e.g., see FIG. 1 and description beginning at page 5, line 1, which shows a computing system 100, application 104, storage device 210A, and replicator 109; FIG. 3 and text beginning page 8, line 17, which depicts memory pool 310, 312). The application is configured to initiate write transactions; and the storage device is configured to store data corresponding to said write transactions. (e.g., page 7, lines 15+). The replicator component configured to: monitor the write transactions; allocate buffers from the memory pool for the write transactions; and automatically modify system resources in response to I/O characteristics of said monitored write transactions. (e.g., see FIG. 5 and page 12, beginning line 18 where the replicator includes a monitor component. The replicator allocates buffer space, page 10, lines 1-2; and modifies resources, FIG. 6 and text beginning page 14, line 1). In modifying resources, the size of the memory pool may be modified (e.g., see page 15, lines 5+). Claim 1 further recites the application, first storage device, and replicator are all within a single node of the system, and system includes a second node with a second storage device coupled to said first node, wherein said replicator component is further configured to convey said write transactions to said second node. (see FIG. 1 and accompanying text beginning at page 5, lien 1).

Claim 9 recites a method which includes initiating write transactions in a first node of a computing system (e.g., FIG. 6, block 602, page 14, line 1+; and FIG. 1 depicting node 120 and system 100). The method further includes allocating buffers from a memory pool for said write transactions (the replicator allocates buffer space, page 10, lines 1-2), and storing data corresponding to said write transactions in a first storage device of the first node (e.g., page 7, lines 15+). Claim 9 also recites monitoring said

write transactions (FIG. 5 and page 12, beginning line 18 where the replicator includes a monitor component); automatically modifying system resources in response to I/O characteristics of said monitored write transactions, wherein said modifying includes modifying a size of said memory pool (e.g., replicator allocates buffer space, page 10, lines 1-2; and modifies resources, FIG. 6 and text beginning page 14, line 1). Claim 9 further recites conveying the write transactions to a second node of the computing system and storing data corresponding to said write transactions in the second node (e.g., see page 5, line 24).

Claim 17 recites a storage medium comprising program instructions. The program instructions are executable to initiate write transactions in a first node of a computing system (e.g., FIG. 6, block 602, page 14, line 1+; and FIG. 1 depicting node 120 and system 100). Claim 17 further recites the instructions are executable to allocate buffers from a memory pool for said write transactions (the replicator allocates buffer space, page 10, lines 1-2), and store data corresponding to said write transactions in a first storage device of the first node (e.g., page 7, lines 15+). Claim 17 also recites the instructions monitor the write transactions (FIG. 5 and page 12, beginning line 18 where the replicator includes a monitor component); automatically modifying system resources in response to I/O characteristics of said monitored write transactions, wherein said modifying includes modifying a size of said memory pool (e.g., replicator allocates buffer space, page 10, lines 1-2; and modifies resources, FIG. 6 and text beginning page 14, line 1). Claim 17 further recites the instructions convey the write transactions to a second node of the computing system and storing data corresponding to said write transactions in the second node (e.g., see page 5, line 24).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1, 2, 4, 6-10, 12, 14-18, and 20 are rejected under 35 U.S.C. § 103(a) as being patentable over U.S. Patent No. 6,625,623 (hereinafter "Midgley") in view of U.S. Patent No. 5,680,573 (hereinafter "Rubin").

2. Claims 21-23 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Midgley in view of Rubin, and further in view of U.S. Patent No. 6,922,791 (hereinafter "Mashayekhi").

VII. ARGUMENT

A. Claims 1, 2, 4, 6-10, 12, 14-18, and 20

The Examiner rejected claims 1, 2, 4, 6-10, 12, 14-18, and 20 as being obvious over Midgley in view of Rubin under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks.

Claim 1 recites a computing system comprising:

> "an application configured to initiate write transactions;
> a first storage device configured to store data corresponding to said write transactions;
> a memory pool; and
> a replicator component configured to:
> monitor said write transactions;
> allocate buffers from said memory pool for said write transactions; and
> automatically modify system resources in response to I/O characteristics of said monitored write transactions, wherein modifying said system resources includes modifying a size of said memory pool;
> wherein said application, first storage device, and replicator are within a first node of said system, and wherein said system includes a second node with a second storage device coupled to said first node, wherein said replicator component is further configured to convey said write transactions to said second node."

As clearly seen from the above claim language, there are two nodes recited in the claim. The first node includes an application, first storage device configured to store write transaction data, and a replicator component. The replicator component is configured to monitor the write transactions, allocate buffers from a memory pool, and automatically modify system resources in response to I/O characteristics of the write transactions. In addition, the replicator component is configured to convey said write transactions to the second node.

In paragraph 3 of the Office Action December 27, 2006, the Examiner asserts that Midgley teaches:

> "a replicator component (see lines 10-52 of column 19) configured to monitor said write transactions (see lines 10-52 of column 19); and automatically modify system resources in response to I/O characteristics of said monitored write transactions … wherein said application, first storage device, and replicator are within a first node of said system and wherein said system includes a second node with a second storage device coupled to said first node, wherein said replicator component is further configured to convey said write transactions to said second node."

However, Applicant disagrees. Midgley discloses (e.g., see Fig. 1) a computer network system 10 which includes a backup server 12 coupled to long term storage 14 (tape library). For example, Midgley discloses:

> "The depicted backup server 12 may also be a conventional workstation system such as a Sun Sparc workstation running a version of the UNIX operating system, or a PC compatible work station running the windows operating system or any other suitable platform." (Midgley, col. 9, lines 43-47).

> "The depicted tape library 14 may be a conventional tape library system of the type commonly employed for backing up data on a computer network. In one particular embodiment, the tape library system 14 is a blank tape library system manufactured by the Quantum corp. of Milpitas, Calif. However, it will be apparent to those of ordinary skill in the art that other tape library systems may be employed without departing from the scope of the invention." (Midgley, col. 8, line 63 – col. 9, line 4).

In addition to the above, Midgley discloses that the backup server 12 includes a number of distinct processes, including a replication process 40, a catalog process 42, and a bandwidth control process 44. For example, Midgley discloses:

> "[T]he synchronization replication process 40 operates in cooperation with the agent processes 30 [located on file servers 18, 20, 22] to create a replica of selected files maintained within the data bases 32, 34 and 38." (Midgley, col. 7, lines 58-61).

> "The catalog process 42 can be a conventional computer process operating on the backup server 12 to collect information form (*sic*) the synchronization and dynamic replication processes to create a database of

the different versions of the target files being stored." (Midgley, col. 10, lines 50-55).

"[T]he back up server 12 may provide a bandwidth control process 44 that may be accessed through the console 24. In the depicted embodiment, the bandwidth control process is shown as operating on the back up server 12, however it will be apparent to those of skill in the art that the bandwidth control process 44 may be located on the data servers 18, 20, 22, or on both the data servers 18, 20, 22 and the back up server 12. The user may employ this process 44 to set a network consumption limit for each backup policy and restore operation."

Accordingly, Midgley discloses a bandwidth control process 44 that may be used to affect resources (network consumption) during backup and restore operations between file servers 18, 20, 22 and long term storage 14. A replication process 40 included in the backup server 12 performs backup and restore operations between the long term storage 14 and the file servers. The agents 30 located on the file servers monitor file accesses to databases 32, 34, 38.

In the above comment, the examiner states Midgley discloses that the application, first storage device, and replicator are within a first node of said system and wherein said system includes a second node with a second storage device coupled to said first node, wherein said replicator component is further configured to convey said write transactions to said second node. However, the portion of Midgley cited by the examiner clearly discloses the operation of a bandwidth control process 44. More specifically, Midgley discloses:

"[T]he back up server 12 may provide a bandwidth control process 44 that may be accessed through the console 24. In the depicted embodiment, the bandwidth control process is shown as operating on the back up server 12, however it will be apparent to those of skill in the art that the bandwidth control process 44 may be located on the data servers 18, 20, 22, or on both the data servers 18, 20, 22 and the back up server 12. The user may employ this process 44 to set a network consumption limit for each backup policy and restore operation. When setting this option, the user may select the bandwidth that is available between the source and backup systems, and specify a consumption limit to be allocated to the synchronization

> and/or dynamic replication processes. If multiple network links are available between the systems, the user may specify the slowest link. Further, the bandwidth control process 44 may include a process for determining, either dynamically, or historically, the available network resources, including network bandwidth and buffer availability, for a given time. The determined resources may be provided by the user through the console process 24, or automatically employed by the bandwidth control process 44 for selecting network consumption limits. . . .
>
> Once the consumption limit is set, the bandwidth control process 44 may throttle the bandwidth usage of the agents 30, synchronization replication process 40 or any replication process by limiting the amount of data to be placed on the network 10 per unit of time. To this end, the bandwidth control process 44 may calculate the bandwidth usage limit based on the maximum percentage of bandwidth the user selected for the operation and the type of network specified. Optionally, the user may vary the network bandwidth consumption for a particular policy over the course of a week. Thus a user could choose to limit consumption during the working hours and allow unlimited consumption at other times." (Midgley, col. 19, lines 10-60)

As seen from the above, the bandwidth control process 44 is a process which "may be located on the data servers 18, 20, 22, or on both the data servers 18, 20, 22 and the back up server 12." Additionally, as clearly described above, and as shown in FIG. 1 of Midgley, the bandwidth control process 44 is distinct from the disclosed replication process 40. As noted above, "the bandwidth control process 44 may throttle the bandwidth usage of the agents 30, synchronization replication process 40 or any replication process by limiting the amount of data to be placed on the network 10 per unit of time." Therefore, the bandwidth control process 44 and replication process 40 are distinct entities which perform different functions.

Claim 1 recites a replicator component configured to "monitor said write transactions; allocate buffers from said memory pool for said write transactions; and automatically modify system resources in response to I/O characteristics of said monitored write transactions." Midgley's bandwidth control process 44 does not perform the function of replication and is not equivalent to the recited replicator component.

In response to Applicant's argument that the bandwidth control process 44 does not perform replication, the examiner states in the Advisory Action that:

> "The bandwidth control process is part of the backup server . . . the backup server has a process that carries out the replication."

However, the fact that the bandwidth control process is located on a server which includes other processes which perform other functions does not thereby mean the bandwidth control process performs those functions.

Additionally, the bandwidth control process 44 is not "configured to convey said write transactions to said second node," as recited in claim 1. Instead, Midgley teaches a different component that performs replication located not on the data servers or clients, but on a backup server, namely, a "synchronization replication process 40" which may *receive* write transaction data from file servers 18, 20, 22. Therefore, not even the replication process of Midgley is configured to "convey said write transactions to said second node" from the first node.

In addition, Applicant finds no teaching or suggestion in Midgley of placing the synchronization replication process 40 anywhere but on the backup server, despite the contrary assertion by the Examiner in paragraph 12 of the office action that "Midgley further discloses the replicator can be within any of the servers 18-22." In the Advisory Action, the examiner states:

> "At lines 25-30 of column 19, Midgley discloses the process can be located on any server in the system."

However, upon review of the cited disclosure it is clear that "the process" is not the replication process 40. Rather, "the process" mentioned above is the bandwidth control process 44. The cited disclosure is as follows:

> "To this end, the back up server 12 may provide a bandwidth control process 44 that may be accessed through the console 24. In the depicted embodiment, the bandwidth control process is shown as operating on the back up server 12, however it will be apparent to

those of skill in the art that the bandwidth control process 44 may be located on the data servers 18, 20, 22, or on both the data servers 18, 20, 22 and the back up server 12."

Therefore, Midgley does not disclose the replication process may be located on any of the servers as suggested by the examiner. Thus, neither does the Applicant find any teaching or suggestion in Midgley that "said application, first storage device, and replicator are within a first node of said system" as is recited in claim 1.

Accordingly, Applicant submits not all of the features of claim 1 are disclosed by the combination of cited art, and claim 1 is patentably distinct for at least the above reasons. As each of independent claims 9 and 17 include features similar to that of claim 1, each of these claims are patentably distinct for similar reasons.

B. Claim 21-23

The Examiner rejected claims 21-23 as being obvious over Midgley and Rubin in view of Mashayekhi under 35 U.S.C. § 103(a). Appellants respectfully traverse this rejection in light of the following remarks. The rejections of claims 21-23 depend upon Midgley as discussed above and are distinguishable from the cited art for at least the above reasons.

CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-2, 4, 6-10, 12, 14-18, and 20-23 was erroneous, and reversal of his decision is respectfully requested.

Respectfully submitted,

/ Rory D. Rankin /
Name Rory D. Rankin
Reg. No. 47,884
ATTORNEY FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: August 6, 2007

VIII. APPENDIX

The claims on appeal are as follows.

1. (Previously presented) A computing system comprising:
   an application configured to initiate write transactions;
   a first storage device configured to store data corresponding to said write transactions;
   a memory pool; and
   a replicator component configured to:
      monitor said write transactions;
      allocate buffers from said memory pool for said write transactions; and
      automatically modify system resources in response to I/O characteristics of said monitored write transactions, wherein modifying said system resources includes modifying a size of said memory pool;
   wherein said application, first storage device, and replicator are within a first node of said system, and wherein said system includes a second node with a second storage device coupled to said first node, wherein said replicator component is further configured to convey said write transactions to said second node.

2. (Original) The computing system as recited in claim 1, wherein said replicator is further configured to record data indicative of said characteristics.

3. (Cancelled).

4. (Previously presented) The computing system as recited in claim 1, further comprising a log volume, and wherein said replicator is further configured to store said write transactions in said log volume.

5. (Cancelled).

6. (Previously presented) The computing system as recited in claim 1, wherein said second node includes a pool of buffers, each of which is configured to store a write transaction received from said first node, and wherein said replicator component is further configured to modify a size of said pool of buffers in said second node in response to said characteristics.

7. (Previously Presented) The computing system as recited in claim 2, wherein said replicator is further configured to:
provide said recorded characteristics for display;
provide guidelines to a user for modifying resources of said system; and
modify said resources based upon user input.

8. (Original) The computing system as recited in claim 6, wherein said replicator component is configured to access said recorded data responsive to detecting an event.

9. (Previously presented) A method comprising:
initiating write transactions in a first node of a computing system;
allocating buffers from a memory pool for said write transactions;
storing data corresponding to said write transactions in a first storage device of the first node;
monitoring said write transactions;
automatically modifying system resources in response to I/O characteristics of said monitored write transactions, wherein said modifying includes modifying a size of said memory pool; and
conveying said write transactions to a second node of the computing system and storing data corresponding to said write transactions in the second node.

10. (Original) The method as recited in claim 9, further comprising recording data indicative of said characteristics.

11. (Cancelled).

12. (Previously presented) The method as recited in claim 9, further comprising storing said write transactions in a log volume.

13. (Cancelled).

14. (Previously presented) The method as recited in claim 9, wherein said second node includes a pool of buffers, each of which is configured to store a write transaction received from said first node, and wherein said method further comprises modifying a size of said pool of buffers in said second node in response to said characteristics.

15. (Previously Presented) The method as recited in claim 10, further comprising:
providing said recorded statistics for display;
providing guidelines to a user for modifying resources of said system; and
modifying said resources based upon user input.

16. (Original) The method as recited in claim 14, further comprising accessing said recorded data responsive to detecting an event.

17. (Previously presented) A machine readable storage medium comprising program instructions, wherein said program instructions are executable to:
initiate write transactions in a first node of a computing system;
allocate buffers from a memory pool for said write transactions;

store data corresponding to said write transactions in a first storage device of the first node;

monitor said write transactions;

automatically modify system resources in response to I/O characteristics of said monitored write transactions, wherein modifying said system resources includes modifying a size of said memory pool; and

convey said write transactions to a second node of the computing system and store data corresponding to said write transactions in the second node.

18. (Previously presented) The storage medium as recited in claim 17, wherein said program instructions are further executable to record data indicative of said characteristics.

19. (Cancelled).

20. (Previously presented) The storage medium as recited in claim 17, wherein said program instructions are further executable to:

convey said write transactions from a first node to a buffer allocated from a pool of buffers within a second node; and

modify a size of said pool of buffers in said second node in response to said characteristics.

21. (Previously presented) The system of claim 1, wherein said second node is configured to serve as a failover node if the first node fails.

22. (Previously presented) The method of claim 9, wherein said second node is configured to serve as a failover node if the first node fails.

23. (Previously presented) The storage medium of claim 17, wherein said second node is configured to serve as a failover node if the first node fails.

IX. EVIDENCE APPENDIX

No evidence submitted under 37 C.F.R. §§ 1.130, 1.131, or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

X. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.